MATERIALIZED VIEWS

Techniques, Implementations, and Applications

Edited by

Ashish Gupta

Inderpal Singh Mumick

Foreword by Jeffrey D. Ullman

The MIT Press
Cambridge, Massachusetts
London, England

Contents

Foreword
Jeffrey D. Ullman xvii

Preface xix

PART I REJUVENATION OF MATERIALIZED VIEWS 1

1 INTRODUCTION TO VIEWS
Ashish Gupta, Inderpal Singh Mumick 3
1 Introduction 3
2 Example 4
3 Views and Integrity Constraints 5
4 Relationship to Rules and Triggers in Active Databases 6
5 Active Elements in Databases 7
6 Historical Work 7

2 MAINTENANCE POLICIES
Ashish Gupta, Inderpal Singh Mumick 9
1 Maintenance Steps 9
2 Maintenance Policies 9
3 Choice of Maintenance Policies 11

3 APPLICATIONS OF MATERIALIZED VIEWS
Ashish Gupta, Inderpal Singh Mumick 13
1 Introduction 13
2 Fast Access, Lower CPU and Disk Load 14
3 Data Warehousing 14
4 Data Replication 23
5 Data Visualization 24
6 Mobile Systems 25

7 Switching Software 26
8 Network Management 27
9 Chronicle Systems in Banking and Billing 28
10 Distributed CD-ROM Services 30
11 Advanced Messaging Services 31
12 Data Pumping 33
13 Query Optimization 33
14 Integrity Constraint Checking 35
15 Distributed Situation Monitoring 36

4 CHALLENGES IN SUPPORTING MATERIALIZED VIEWS
Ashish Gupta, Inderpal Singh Mumick 39
1 Introduction 39
2 General Problems 39
3 Application-Specific Problems 42
4 Database Systems 47

PART II APPLICATIONS OF MATERIALIZED VIEWS 49

5 DATA INTEGRATION USING SELF-MAINTAINABLE VIEWS
Ashish Gupta, H. V. Jagadish, Inderpal Singh Mumick 53
1 Introduction 53
2 Background and Notation 54
3 Self-Maintenance for SPJ Views 55
4 Related Work 57
5 Conclusions and Future Work 57

6 USING OBJECT MATCHING AND MATERIALIZATION TO INTEGRATE HETEROGENEOUS DATABASES
Gang Zhou, Richard Hull, Roger King, Jean-Claude Franchitti 59
1 Introduction 59
2 Preliminaries 61
3 An Example Problem and Its Solution 63
4 A Taxonomy of the Solution Space for Data Integration 67
5 Towards a General Tool for Building Integration Mediators 70
6 Current Status and Future Research 75

7 OPTIMIZING QUERIES WITH MATERIALIZED VIEWS
Surajit Chaudhuri, Ravi Krishnamurthy, Spyros Potamianos, Kyuseok Shim 77
1 Introduction 77

2 Overview of Our Approach 79
3 Equivalent Queries: Generation of MapTable 80
4 Join Enumeration 86
5 Discussion 90
6 Related Work 91
7 Summary 92

8 ANSWERING QUERIES USING VIEWS
Alon Y. Levy, Alberto O. Mendelzon, Yehoshua Sagiv, Divesh Srivastava 93
1 Introduction 93
2 Preliminaries 93
3 Complexity of Finding Rewritings 96
4 Finding Redundant Literals in the Rewritten Query 101
5 Related Work 106

9 ADAPTING MATERIALIZED VIEWS AFTER REDEFINITIONS
Ashish Gupta, Inderpal Singh Mumick, Kenneth A. Ross 107
1 Introduction 107
2 The System Model 110
3 `SELECT-FROM-WHERE` Views 112
4 Aggregation Views 117
5 Union and Difference Views 121
6 Multiple Changes to a View Definition 122
7 Conclusions 123

10 THE DYNAMIC HOMEFINDER: EVALUATING DYNAMIC QUERIES IN A REAL-ESTATE INFORMATION EXPLORATION SYSTEM
Christopher Williamson, Ben Shneiderman 125
1 Introduction and Background 125
2 Dynamic Queries Interface to Real-Estate 127
3 User Experiment 128
4 Experiment Results 130
5 Discussion and Conclusions 131
6 Acknowledgments 133

PART III MAINTENANCE OF MATERIALIZED VIEWS 141

11 MAINTENANCE OF MATERIALIZED VIEWS: PROBLEMS, TECHNIQUES, AND APPLICATIONS
Ashish Gupta, Inderpal Singh Mumick 145
1 Introduction 145
2 The Idea Behind View Maintenance 149
3 Using Full Information 149
4 Using Partial Information 153
5 Open Problems 156

12 INCREMENTAL MAINTENANCE OF RECURSIVE VIEWS: A SURVEY
Guozhu Dong 159
1 Dimensions of View Maintenance Methods 159
2 Using Relational Queries as Maintenance Algorithms 160
3 Recursive Maintenance Methods 162

13 EFFICIENTLY UPDATING MATERIALIZED VIEWS
José A. Blakeley, Per-Åke Larson, Frank Wm. Tompa 163
1 Introduction 163
2 Previous Work 164
3 Notation and Terminology 165
4 Relevant and Irrelevant Updates 165
5 Differential Re-evaluation of Views 169
6 Conclusions 175

14 MAINTAINING VIEWS INCREMENTALLY
Ashish Gupta, Inderpal Singh Mumick, V.S. Subrahmanian 177
1 Introduction 177
2 Related Work 179
3 Notation 180
4 Incremental Maintenance of Nonrecursive Views Using Counting 181
5 Implementation Issues and Optimizations 183
6 Negation and Aggregation 185
7 Incremental Maintenance of Recursive Views 188
8 Conclusions and Future Work 188

15 INCREMENTAL MAINTENANCE OF VIEWS WITH DUPLICATES
Timothy Griffin, Leonid Libkin 191
1 Introduction 191
2 Basic Notation 194
3 Preliminaries 197
4 Change Propagation Algorithm 200
5 Top-Level Aggregate Functions 202
6 Complexity Analysis 203
7 Related Work 205
8 Further Work 207

16 ALGORITHMS FOR DEFERRED VIEW MAINTENANCE
Latha Colby, Timothy Griffin, Leonid Libkin, Inderpal Singh Mumick, Howard Trickey 209
1 Introduction 209
2 Preliminaries 213
3 View Maintenance Scenarios 216
4 Exploiting Duality 219
5 Algorithms and Policies 223
6 Related Work 226
7 Future Work 226

17 INCREMENTAL EVALUATION OF DATALOG QUERIES
Guozhu Dong, Rodney Topor 229
1 Introduction 229
2 Basic Concepts 230
3 Regular Chain Programs 232
4 Arbitrary Datalog Programs 237
5 Conclusions and Discussion 238

18 VIEW MAINTENANCE ISSUES FOR THE CHRONICLE DATA MODEL
H. V. Jagadish, Inderpal Singh Mumick, Abraham Silberschatz 241
1 Introduction 241
2 The Chronicle Data Model 243
3 Complexity of a Chronicle Model 245
4 Summarized Chronicle Algebra 246
5 Issues in the Chronicle Model 249
6 Related Work 250
7 Conclusions 251

19 VIEW MAINTENANCE IN A WAREHOUSING ENVIRONMENT
Yue Zhuge, Hector Garcia-Molina, Joachim Hammer, Jennifer Widom 253
1 Introduction 253
2 Related Research 258
3 Correctness 258
4 Views and Queries 260
5 The ECA Algorithm 262
6 Performance Evaluation 268
7 Conclusion 272

20 EFFICIENT MAINTENANCE OF MATERIALIZED MEDIATED VIEWS
James Lu, Guido Moerkotte, Joachim Schü, V.S. Subrahmanian 275
1 Introduction 275
2 Preliminaries 277
3 Updating Views 281
4 Maintaining Views when External Changes Occur 290
5 Discussion 292
6 Conclusion and Future Work 292

21 UPDATING DERIVED RELATIONS: DETECTING IRRELEVANT AND AUTONOMOUSLY COMPUTABLE UPDATES
José A. Blakeley, Neil Coburn, Per-Åke Larson 295
1 Introduction 295
2 Notation and Basic Assumptions 297
3 Irrelevant Updates 299
4 Autonomously Computable Updates 304
5 Discussion 319

22 QUERIES INDEPENDENT OF UPDATES
Alon Y. Levy, Yehoshua Sagiv 323
1 Introduction 323
2 Preliminaries 324
3 Detecting Independence 327
4 Uniform Equivalence 332
5 Concluding Remarks 338
6 Acknowledgments 338

PART IV MATERIALIZED VIEWS AND OLAP 339

23 IMPLEMENTING DATA CUBES EFFICIENTLY
Venky Harinarayan, Anand Rajaraman, Jeffrey D. Ullman 343
1 Introduction 343
2 The Lattice Framework 348
3 The Cost Model 350
4 Optimizing Data-Cube Lattices 352
5 The Hypercube Lattice 358
6 Conclusions and Future Work 359

24 ON THE COMPUTATION OF MULTIDIMENSIONAL AGGREGATES
Sameet Agarwal, Rakesh Agarwal, Prasad M. Deshpande, Ashish Gupta, Jeffrey F. Naughton, Raghu Ramakrishnan, Sunita Sarawagi 361
1 Introduction 361
2 Optimizations Possible 363
3 Sort-Based Methods 364
4 Hash-Based Methods 367
5 Experimental Evaluation 369
6 Contributions of this Part 372
7 Options for Computing the CUBE 372
8 The Overlap Method 373
9 Implementation and Results 377
10 Conclusions and Summary 379

25 MAINTENANCE OF DATA CUBES AND SUMMARY TABLES IN A WAREHOUSE
Inderpal Singh Mumick, Dallan Quass, Barinderpal Singh Mumick 387
1 Introduction 387
2 Motivating Example 388
3 Background and Notation 392
4 Basic Summary-Delta Maintenance Algorithm 396
5 Efficiently Maintaining Multiple Summary Tables 401
6 Performance 404
7 Related Work and Conclusions 405

PART V IMPLEMENTATION AND PERFORMANCE ANALYSIS OF MATERIALIZED VIEWS 409

26 ORACLE7 SNAPSHOTS
Alan Downing, Ashish Gupta, Robert Jenkins, Harry Sun 413
1 Introduction 413
2 Overview of Snapshots in Oracle7 414
3 System Architecture 414
4 Refresh Algorithm 416
5 Conclusions 420

27 CHECKING INTEGRITY AND MATERIALIZING VIEWS BY UPDATE PROPAGATION IN THE EKS SYSTEM
Laurent Vieille, Petra Bayer, Volker Küchenhoff, Alexandre Lefebvre 421
1 Introduction 421
2 Background 423
3 Basic Notions of Update Propagation 426
4 Testing Policy 428
5 Relevant Propagation Rules 430
6 Managing Propagation Rules 433
7 Discussion and Previous Work 436
8 Conclusion 439

28 DERIVING PRODUCTION RULES FOR INCREMENTAL VIEW MAINTENANCE
Stefano Ceri, Jennifer Widom 441
1 Introduction 441
2 View Definition Language 444
3 Production Rule Language 445
4 Motivation 446
5 Top-Level Table References 447
6 Positively Nested Subqueries 454
7 Negatively Nested Subqueries 457
8 Set Operators 460
9 System Execution 461
10 Conclusions and Future Work 462

29 THE HERACLITUS DBPL WITH APPLICATION TO ACTIVE DATABASES AND DATA INTEGRATION
Gang Zhou, Richard Hull, Shahram Ghandeharizadeh 465
1 Introduction 465

2 Deltas, Virtual States, and Active Database Execution Models 466
3 Heraclitus[Alg,C] 471
4 The Implementation of Heraclitus[Alg,C] 477
5 Heraclitus and Maintenance of Materialized Data 481
6 Conclusions 483

30 AN INCREMENTAL ACCESS METHOD OF VIEW CACHE: CONCEPT, ALGORITHMS, AND COST ANALYSIS
Nicholas Roussopoulos 485
1 Introduction 485
2 Incremental Access Methods for Views 486
3 The IAM Algorithms 497
4 The IAM Cost Model 500
5 ADMS Implementation and Experiments 505
6 Conclusions 508
7 Acknowledgments 509

31 A PERFORMANCE ANALYSIS OF VIEW MATERIALIZATION STRATEGIES[1]
Eric N. Hanson 511
1 Introduction 511
2 View Materialization Strategies 512
3 Performance Comparison 516
4 Conclusion 530

32 JOIN INDEX, MATERIALIZED VIEW, AND HYBRID-HASH JOIN: A PERFORMANCE ANALYSIS
José A. Blakeley, Nancy L. Martin 535
1 Introduction 535
2 Methods 536
3 Performance Analysis 538
4 Results 547
5 Conclusion 549

References 551

Source Notes 578

Contributors 580

Index 587